

(October 13th, 2014)

If you know someone who you think would benefit from being an Insider, feel free to forward this PDF to them so they can sign up [here](#).



Quick Tips for our Insider friends!

Hey Insiders!

This newsletter is coming to you from Chicago where we're just starting IEPTO2 after teaching IEPTO1 and IE0 last week. When we're not teaching we're working on more Pluralsight courses so expect to see 6 or so more before the end of the year.

2015 classes are open for registration! Late last week I opened up our spring 2015 classes in Chicago and Bellevue for registrations. All classes have a discount for registering in 2014 (US\$125 off IE0 and US\$200 off all 5-day classes) so get those training budgets working for you before you lose them at the end of the year! See [here](#) for details.

The most recent book I've read is George Dyson's [Turing's Cathedral](#) (Dyson is the son of [Freeman Dyson](#), the celebrated Professor of Physics). This is a really interesting book that charts the birth of the computer during the 1940s and 1950s in the decade after WWII through the work of the group based at the Institute of Advanced Study in Princeton, NJ. The center of the group (and book) was John von Neumann, not Turing, which makes the book's title curious, but the book also describes the other major engineering contributors. Recommended.

Please [let us know](#) if you liked what you read/saw here and/or have any suggestions for future Quick Tips.

Note: you can get all the prior Insider newsletters [here](#).

Paul's Ponderings

Availability groups (AGs) are pretty cool, and one of the most useful features of them is the ability to read directly from one of the secondary replicas. Before, with database mirroring, the only way to access the mirror database was through the creation of a database snapshot, which only gave a single, static view of the data. Readable secondaries are constantly updated from the primary so are far more versatile as a reporting or non-production querying platform.

But I bet you didn't know that using this feature can cause big performance problems on your primary replica?

As with most things in life, you don't get anything for free. Readable secondaries are really useful, but there is a performance tradeoff you need to be aware of.

All queries that are executed against a readable secondary are automatically run using read-committed snapshot isolation. This means they do not require share locks and so will not block any database changes being replayed from the primary replica (i.e. the constant redo of log records on the secondary replica that have been sent from the primary replica).

To do this requires the use of the versioning system, where (simplistically) pre-change versions of records are copied into the version store in tempdb and queries work out which version of the record is the correct one for them to process, based on the query's starting time. All records that change get a 14-byte tag added on the end of the record that allows a query to see if this is the correct record, and if not to follow a pointer to the previous version of the record in the version store. This has been the mechanism since snapshot isolation and read-committed snapshot isolation were introduced in SQL Server 2005.

Now consider this: all AG replicas are exact copies of the primary replica. So how can versioning work on the readable secondary, adding 14-byte tags to some records? That must break the 'exact copy' rule, right?

Well, yes, it would... if the primary replica didn't also change.

When a readable secondary is configured in an AG environment, all changing records on the primary replica start getting empty 14-byte versioning tags added to them. This is so that the 14-bytes of extra space on the record is noted in the transaction log and replayed on the secondary replicas, allowing the readable secondary to make use of the empty 14-byte space to store the versioning tag it needs.

This doesn't break the 'exact copy' rule because the 14-bytes isn't used for anything to do with recovery, there just has to be 14-bytes there.

So versioning tags start getting added to changing records on the primary (to be clear, it doesn't turn on versioning on the primary) so table and index records start to get 14-bytes longer. And what happens when records get longer on pages where there isn't enough space? Page splits in your indexes (and forwarded records in heaps – but I'll concentrate on indexes here) leading to low page densities (wasted disk space and buffer pool memory), logical fragmentation (poor range scan performance), and a bunch of extra, expensive log record generation from the page splits themselves.

To counteract this, you'll need to implement (and/or possibly lower existing) fillfactors on your indexes and even potentially start doing index maintenance on indexes that may not have required it previously. Quite an insidious problem that can be hard to figure out unless you know what's going on under the covers!

See the following blog posts of mine for more info:

- [Performance issues from wasted buffer pool memory](#)
- [How expensive are page splits in terms of transaction log?](#)
- [How to choose a good index fillfactor?](#)

- [Easy automation of SQL Server database maintenance](#)

Call to action: If you're implementing readable secondaries in your AG configuration, make sure that you also investigate and implement index fillfactors in the database so that the versioning tags that are added under the covers don't start causing page splits and fragmentation.

I'm curious to hear your thoughts on this performance impact of using AGs, so please feel free to [drop me a line](#), always treated confidentially, of course.

Video Demo

As we're all teaching right now, I chose to share a clip from Jon's Pluralsight course [SQL Server: Deadlock Analysis and Prevention](#). In the clip, Jon explains how (very common) reverse object order deadlocks occur and can be diagnosed.

The video is just over 4 minutes long and you can get in WMV format [here](#).

You can get the demo code [here](#).

Enjoy!

SQLskills Offerings

We've released the first batch of classes for 2015! We may add another class in February, and there will likely be some in the second half of 2015, but for now these are all we have on the schedule. All classes have discounts for registering in 2014.

Finally, to help your boss understand the importance of focused, technical training, we've added a few new items to help you justify spending your training dollars with us:

- [Letter to your boss explaining why SQLskills training is worthwhile](#)
- [Community blog posts about our classes](#)
- [Immersion Event FAQ](#)

Upcoming Immersion Events

Sydney, NSW, Australia

- December 8-12, 2014: **IEPTO1**: Immersion Event on Performance Tuning and Optimization – Part 1 (formerly IE1)

Chicago, IL

- April 27-29, 2015: **IE0**: Immersion Event for the Accidental/Junior DBA

- April 27-May 1, 2015: **IEPTO1**: Immersion Event on Performance Tuning and Optimization – Part 1 (formerly IE1)
- May 4-8, 2015: **IEPTO2**: Immersion Event on Performance Tuning and Optimization – Part 2 (formerly IE2)
- May 4-8, 2015: **IEBI**: Immersion Event on Business Intelligence
- May 11-15, 2015: **IEHADR**: Immersion Event on High Availability and Disaster Recovery (formerly IE3)

Bellevue, WA

- June 8-12, 2015: **IEPTO1**: Immersion Event on Performance Tuning and Optimization – Part 1 (formerly IE1)
- June 15-19, 2015: **IEPTO2**: Immersion Event on Performance Tuning and Optimization – Part 2 (formerly IE2)

See [here](#) for the main Immersion Event Calendar page that allows you to drill through to each class for more details and registration links.

Fall SQLintersection

This year our Fall SQLintersection conference will be the week of November 10th in Las Vegas. See [here](#) for details. We've completed our Fall lineup by adding two more fantastic Microsoft speakers.

Shep Sheppard will be joining us. In the Spring, we were so impressed that we asked if he would join us again in the Fall and we're excited to have him for three sessions (**SQL Server Troubleshooting the CSS Way; SQL Server 2014 Gems; and Microsoft SQL Server and the Cloud – Options, Strategies, Benefits**).

The other Microsoft speaker to join us is SQL Server Storage Engine Program Manager, Kevin Farlee. Also a request granted (given his long-time background with SQL), Kevin will be delivering three sessions (**In-Memory OLTP Overview; In-Memory Data Warehousing Overview; and SQL Server I/O – Options, Performance, Benefits**).

Our Fall lineup is one of our best yet; targeting the pain points you're having TODAY with some sessions to help you get ready for tomorrow. We hope you can join us!

Don't forget to use the discount code "SQLskills" (without the quotes and it isn't case-sensitive) and you can save \$50 off registration.

Summary

We hope you've enjoyed this issue - we really enjoy putting these together.

If there is anything else you're interested in, we'd love to hear from you - [drop us a line](#).

Thanks,

Paul and Kimberly

Paul@SQLskills.com and Kimberly@SQLskills.com